arXiv:1612.06749v1 [cs.DC] 20 Dec 2016

GuStL – An Experimental Guarded States Language

Oskar Schirmer

Göttingen, 2016-12-20

Abstract

Programming a parallel computing system that consists of several thousands or even up to a million message passing processing units may ask for a language that supports waiting for and sending messages over hardware channels.

As programs are looked upon as state machines, the language provides syntax to implement a main event driven loop.

The language presented herewith surely will not serve as a generic programming language for any arbitrary task. Its main purpose is to allow for a prototypical implementation of a dynamic software system as a proof of concept.

Contents

1 Introduction **3**

2 Syntax **3**

3 Semantics **5**

4 Process Creation **8**

5 Segmented Program Development **9**

6 Compiler **10**

7 Discussion **10**

1 Introduction

The *Guarded States Language* essentially does combine well known concepts of existing languages, while any concept not needed for its main purpose is avoided. Its basic concept is imperative sequential operations. Syntactically it is inspired by Algol ([1963nb]) and Pascal ([1975jw]). Main program execution is organized as a state machine as with SDL ([2007ir]). States are subject to guards, similar to those introduced in [1975ed]. For channel operations, input and output, the notation follows [1978ch].

A software system may consist of various processes running in parallel, each of which is a program executed sequentially, one instruction at a time. Channels may be set up dynamically to allow transmission of data, i.e. messages, from one process to the other.

Depending on the current process state, different guards may be evaluated, causing some statement sequence to be executed as soon as the condition of the corresponding guard is met.

Statements serve to implement traditional imperative operations, such as assigments, loops, conditional execution and subroutine invocation.

The language does only provide features restricted to the processing unit the program is executed on. As a consequence, and as dynamic process creation does imply activity of various system processes – especially in the case of creating a remote process –, the language does not provide syntactic means to describe instantiation of a new process.

2 Syntax

A program is made up of a sequence of symbols. A symbol may be one of: special symbol, reserved word, identifier, number, or string.

Symbols may be separated by white space, i.e. blank, tabulator, or newline characters. Any character sequence starting with a left brace up to and including the next right brace is considered white space, too, as long as it does not start within a string. The latter allows insertion of comments. Note, that nested comments are not valid.

For the special symbols and reserved words, see the formal syntax definition below. Reserved words are to be written in lower case.

An identifier is composed of a sequence of upper and lower case letters, digits, and the underline character. It must not start with a digit. Identifiers are case sensitive.

Numbers are given in either decimal, hexadecimal, or character notation. For decimal notation a number is written as a sequence of digits, for hexadec-

imal notation, it is preceded by `0x`. A number given in character notation is the characters implicit textual representation[1] enclosed by single quotes. Scientific notation is not supported.

A string is a series of characters – i.e. their implicit textual representations – enclosed by double quotes. To include a double quote in a string, it is duplicated. A string is encoded as an array at one word per character.

A valid sentence for a program must match the following partially formalised syntax description.[2]

process ::= { *constdef* | *wordsdecl* | *proceduredecl* | *functiondecl* }
"`process`" $ident_{!process}$ "`(`" $ident_{!port}$ ["`,`" $ident_{!const}$] "`)`"
[*statesdecl*] [*portsdecl*]
{ *constdef* | *wordsdecl* | *proceduredecl* | *functiondecl* }
"`start`" { *statement* } { *guardedstate* } "`stop`"
statesdecl ::= "`state`" $ident_{!state}$ { "`,`" $ident_{!state}$ }
portsdecl ::= "`port`" $ident_{!port}$ { "`,`" $ident_{!port}$ }
constdef ::= "`const`" ($ident_{!const}$ "`=`" *expression* | $ident_{!constarray}$
"`[`" [*expression*] "`]`" "`=`" *constarray* { "`,`" *constarray* })
constarray ::= *expression* | "`"` *chars* "`"`
wordsdecl ::= "`word`" *worddecl* { "`,`" *worddecl* }
worddecl ::= $ident_{!word}$ | $ident_{!array}$ "`[`" *expression* "`]`"
proceduredecl ::= "`procedure`" $ident_{!procedure}$ "`(`" *formalparameters* "`)`"
{ *wordsdecl* } "`do`" { *statement* } "`return`"
functiondecl ::= "`function`" $ident_{!function}$ "`(`" *formalparameters* "`)`"
{ *wordsdecl* } "`do`" { *statement* } "`return`" *expression*
formalparameters ::= [*formalparameter* { "`,`" *formalparameter* }]
formalparameter ::= "`const`" $ident_{!constarray}$ "`[`" [*expression*] "`]`" |
$ident_{!array}$ "`[`" [*expression*] "`]`" | $ident_{!word}$ | "`port`" $ident_{!port}$
guardedstate ::= "`on`" *statelist* ["\" *guard*] "`:`" { *statement* }
statelist ::= $ident_{?state}$ { "`,`" $ident_{?state}$ }
guard ::= $ident_{?port}$ "`!`" | *reception* | *expiration*
statement ::= *conditional* | *repetition* | *transition* | *assignment* |
transmission | *procedurecall* | *inlineasm*

[1] It is assumed that the usual implicit textual representation for characters is its Unicode code point as defined in `[2009uc]`

[2] For reasons of simplicity, it is given in a somewhat modified version of the notation inspired by Pāṇini (`[1840ob]`) and initially used by Backus (`[1959jb]`). To achieve a fully formalised definition, including semantics, one would need to describe the language by using, e.g., an extended affix grammar (`[1997dw]`). The scope of identifiers is hinted with a subscript $_{!scope}$, where identifiers are declared, and with a subscript $_{?scope}$, where they are assumed to be already declared in the scope given

conditional ::= "`if`" *expression* "`then`" { *statement* }
{ "`elseif`" *expression* "`then`" { *statement* } }
["`else`" { *statement* }] "`done`"
repetition ::= ["`while`" *expression*] "`repeat`" *expression* "`times`"
{ *statement* } ("`done`" | "`until`" *expression*)
transition ::= "`next`" $ident_{?state}$
assignment ::= *variable* "`:=`" (*expression* | *inlineasm*)
transmission ::= $ident_{?port}$ "`!`" (*expression* | "`end`" | "`pause`")
reception ::= $ident_{?port}$ "`?`" (*variable* | "`end`")
variable ::= $ident_{?word}$ | $ident_{?array}$ "`[`" *expression* "`]`" | $ident_{?port}$
expression ::= *simpleexpr*
[("`=`" | "`<>`" | ("`<`" | "`>`" | "`<=`" | "`>=`") ["`$`"]) *simpleexpr*]
simpleexpr ::= *term* { ("`+`" | "`-`" | "`|`" | "`^`" | "`or`") *term* }
term ::= *factor*
{ ("`*`" | "`&`" | "`and`" | "`<<`" | ("`/`" | "`%`" | "`>>`") ["`$`"]) *factor* }
factor ::= [*unaryop*] (*variable* | *constant* | *functioncall* | "`now`"
| *reception* | "`(`" *expression* "`)`")
unaryop ::= "`-`" | "`~`" | "`not`"
constant ::= *number* | "`'`" *char* "`'`" | "`#`" ($ident_{?array}$ | $ident_{?constarray}$) |
$ident_{?const}$ | $ident_{?constarray}$ "`[`" *expression* "`]`"
procedurecall ::= $ident_{?procedure}$ "`(`" *actualparameters* "`)`"
functioncall ::= $ident_{?function}$ "`(`" *actualparameters* "`)`"
actualparameters ::= [*actualparameter* { "`,`" *actualparameter* }]
actualparameter ::= *expression* |
($ident_{?array}$ | $ident_{?constarray}$) ["`[`" *expression* "`]`"]
expiration ::= "`after`" *expression*
inlineasm ::= "`asm`" *singleasm* { "`,`" *singleasm* }
singleasm ::= [*unaryop*] *constant* | "`(`" *expression* "`)`"

3 Semantics

The process is a state machine, based upon an arbitrary number of states. As a special case, the number of states may be zero. In any other case, all state identifiers must be declared next to the process header.

To allow external communication, a process may want to transmit or receive data over channels. To set up a channel at runtime, it must be dynamically assigned to a port, which in turn is static. Any port is referred to by an identifier declared before being used.

Constant identifiers may be declared to stand for constant words or con-

stant arrays. The value of a constant array is given as a list of numbers, where all or parts of the list may be replaced by a string.

The basic unit of data to operate on is a word. It represents a natural or integer number, which depends on which operation is performed on it. Different basic types to restrict specific operations or to determine details of an operation – like signedness – do not exist.[3] A variables identifier may either be declared to refer to a single word, or, when declared as an array, to refer to a sequence of words. Variables must be declared prior to being used.

Subroutines – both functions and procedures – may be declared to implement sequences of statements for multiple use in the program. Upon termination, functions provide a word return value to be used in the invocating expression, while procedures do not return any value. Subroutines may be declared to take a list of parameters, which may be words or arrays or ports. Furthermore, subroutines allow declaring local single word variables, which are instantiated at subroutine invocation and cease at the end of the subroutines execution. Subroutines must be declared prior to being used.

Statements in the program or in a subroutine are executed sequentially, one by one. For a conditional statement, a series of expressions is evaluated, until one is met – i.e. the expression evaluates to non-zero – and the corresponding list of statements is executed. If no condition is met and an alternative branch is given with `else`, that alternative branches statements are executed. A repetition statement is executed for a maximum number of iterations. Furthermore, it is not iterated or not executed at all, when the initial condition is not met, and it is no longer iterated when the terminating condition is met.

The right side of an assignment is evaluated and the result is assigned to the word or port on the left side. A value assigned to a port takes it for a destination port number and prepares a channel to be set up from the local to the destination port.

For channel data transmission, the port to communicate over is given on the left side of the statement. The data to transmit is given on the right side, this may either be an arbitrary value, or a special token to end or pause a message.

Channel data reception is not a statement, but a factor in an expression. As with transmission, the port to receive data from is given on the left side, while on the right side is either a variable to store data into or the end symbol. Both variants consume an incoming end token when it is available,

[3]Various programming languages define signedness as a property of each single data unit. This way, wherever data units of different type are subject to an operation, implicit conversion rules need to be applied, an everlasting source of confusion. To avoid this mess, GuStL makes signedness be a property of the operation itself.

but if data is pending, it is only consumed in the first case and stored into the variable. Same as with the assignment, storing the newly received word to a port sets the ports destination. The factor evaluates to non-zero whenever the expected token is received, which is data for the first variant and an end token for the second. Otherwise, it evaluates to zero.

Calling a subroutine implies evaluation of the list of actual parameters, which must match the list of formal parameters given in the subroutines declaration. For formal word parameters, the corresponding expression is evaluated and its result used in the subroutine (call by value). For formal array parameters, the reference to the corresponding variable is determined and used in the subroutine (call by reference). For formal port parameters, the given port identification is used in the subroutine (call by reference). For a subroutine with more than one parameter the order of parameter evaluation is not defined.

The transition statement is used to change the machines state. It must not be used in a subroutine. Additionally, every execution path in the main program must end in a transition statement.

The main program starts with a list of statements for initialisation purposes, and ends with a list of statements to handle proper process termination. In between, an arbitrary number of guarded statements is provided. Whenever the initialising list of statements or another guarded statement of the program has finished executing the next guarded statement may be selected for execution. Only those guarded statements, that match the current state, are considered. From these, one random guarded statement is selected, for which the condition of its guard is met. When no condition of any guard is met at all, the process is stalled until at least one condition is met.

A guard does either check a port for transmission readiness, for availability of data to receive, for an end token received on a port, or for a given time being expired. When data is available and the corresponding guarded statement is selected for execution, the data received is actually stored into the variable given. If for a state a guard is declared to check for availability of input data, another guard to check for an end token on the same port is obligatory for the same state. For one state several different guards may be declared, provided they do not check for the same condition on the same port, and that no two guards check for expiration.

By default, when evaluating an expression, all operations are performed as unsigned. For various operators, however, appending a `$` sign to the operator itself may be used to ask for signed operation. Signed division is defined to be Euclidean (see e.g. `[2001dl]`). Any expression is composed from one or two simple expressions, simple expressions are composed from terms, terms are composed from factors. A factor may be a single word or a word from

an array, a constant number or an element from a constant array, the size of an array, the call of a function – using its return value, the current time (`now`), the result of checking an input port for having received an end token, or another expression. Any factor may be prepended by a unary operator. The order in which parts of an expression are evaluated is not defined.

To allow machine specific operations, inline assembly is allowed both as a statement and as the right side of an assignment. Inline assembly is given as a list of opcodes, but may include also ordinary expressions to be evaluated. Opcodes depend on the implementation, see section 6. An expression included in inline assembly will leave the corresponding result on the data stack.

4 Process Creation

For traditional single node computers according to the von-Neumann architecture (`[1945jn]`), the basic resource to be allocated by processes at runtime is memory, taken from a uniform pool of contiguous memory words.

For machines made up of large amounts of parallel processing units, each with its own local memory, the allocatable quantity is a single processing unit. Thus, allocating resource does no longer mean to reserve a portion of memory, but to start a process on a previously idle processing unit: *process creation*.

The program to run on the allocated processing unit has to be identified, e.g. by name, which in turn may be variable, and does not refer to data held inside the current program, so for the language to natively support process creation a hypothetical rule could look like the following – where the identifier would refer to a string denoting the new processes name:

creation ::= "`new`" $ident_{?array}$ "`(`" *expression* "`)`"

However, depending on the hardware, starting a process may highly depend on supportive software, i.e. the operating system. The compiler would need to produce operating system dependent code for process creation. To avoid complexity and increase flexibility, means for process creation shall instead be provided by the operating system as an ordinary subroutine. An outline on how it could be declared is given hereafter.

Once the new process is being executed, the process that started it may need to communicate with it. To do so, it needs to know at least one port of the new process to attach to, the *control port*. The subroutine for process creation will return its global identification, so it must be a function. For the new process, the port is declared as the first parameter to the process

declaration header.

As the processing unit and its local memory are – once allocated – an indivisible unit, the amount of memory available during lifetime of the process cannot be changed. Only at the time of creation of the process, the amount of memory given to it may be chosen. For this purpose, the subroutine to create the new process not only needs to know its name, but also a number, which will be available inside the new process as a constant, the *dimension*. To make it usable in a running program, it is declared to be a constant as the second parameter to the process header. It is optional and defaults to zero.

Hardware driver processes may need to access specific hardware locations, so it is crucial to start such a process on a given processing unit. For other reasons, it may be desirable to choose the processing unit from a restricted subset of all units available, so some flags might possibly be helpful. An additional parameter to the function starting the process is used to commit the information needed. The function may be declared as follows:

```
function new(name[], dimension, extra)
```

5 Segmented Program Development

When writing a new program, parts of it may be similar or even identical to portions of a previously implemented program. The wish to reuse these portions and avoid repeated work should represent one of the basic desires of every programmer.

The most simple solution would be to simply *copy and paste* the portion in question to the new program. While in earlier days this was frowned upon for the reason of wasting storage, today the untraceable propagation of programming mistakes is a major concern.

Separate compilation units had the advantage to circumvent the need to recompile all of the code used in a program by keeping the program divided into portions even in its compiled state. This was useful in times of slow compiling machines, but bears the risk of binary incompatibilities when interfaces are changed in source files. Further, an additional *linker* program or a linking loader is needed to enable program execution.

Assembly of the program at compile time may be achieved using *include files*, a method to splice a set of source files.[4] But it may also lead to confusion of which include file is used where, especially when using nested inclusions.

[4]Using C ([1978kr]) usually involves a combination of both linkable program portions, called *libraries*, and include files for interface declaration

To avoid the latter problem, source file composition shall be restricted to simple concatenation. In fact, neither the language nor the compiler need to take any provisions for it, as the concatenation may well be performed by the invoking entity before feeding the complete source program into the compiler itself.

The fact that insertion of one source file into another is impossible makes it necessary to allow declarations before the process declaration of a program, as generic declarations need to be prepended to the main program.

6 Compiler

A first compiler for GuStL has been implemented to produce code for the experimental machine *NOP* (`[2016os]`). The compiler itself is implemented in C, running in a Unix environment. Its single pass design follows the basic concept of `[1977nw]`. It compiles source from *standard input* to binary code on *standard output*. Matching the target machine, words are implemented as 32 bit wide, least significant byte first, and signed numbers are represented in two's complement. The implicit textual representation for characters is its Unicode code point as defined in `[2009uc]`, encoded using UTF-8 in the source file (see also `[2003fy]`).

The binary code itself is preceded by a six word header to allow direct use in an operating system environment:

magic	`0x85cf80cf`
flags	`0`
d_1	data memory size: $\mathrm{ld}_1\text{-}\mathrm{ld}_0 = d_1 \cdot \mathit{dimension} + d_0$, for *dimension* see page 9, see also `[2016os]`
d_0	
code	code and constants memory size: $\mathrm{lc}_1\text{-}\mathrm{lc}_0$, see `[2016os]`
entry	word offset to initial instruction, relative to lc_0, see `[2016os]`

7 Discussion

The main purpose of the programming language GuStL is to allow for a prototypical implementation of a dynamic software system, and its design is restricted to features that are needed to serve this purpose.

There is only very limited support for string handling, namely definition of constant arrays. It might be useful to provide more support for string handling, e.g. transmission and reception of sequences instead of a single word. However, the latter might reduce program flow determinism, because time behavior for the processing of a sequence potentially is more complex

than for a single word, scattering non-determinism from controlling guards into imperative program flow.

Direct recursion in a subroutine is supported, but recursion involving several subroutines is not, as there is no forward declaration. With local memory being quite limited, it is arguable whether system abstraction is adequate to justify support of massive recursion.

Likewise, there is no specific support for floating point calculations, list processing, bit field computation or set arithmetics[5], data pointers or complex data structures. Where this is perceived as a deficiency, it might be worth considering the choice or design of a different language that better fits the particular project context.

Execution of code in parallel may be handled on different levels of a system. GuStL is designed to implement a dynamic software system, and thus it provides parallel execution on process level, where every process start requires to dynamically determine a free resource, the loading of code, and setup of control communication. For static software systems, e.g. embedded control, things are different in that all parts of the software needed at runtime are known prior to system start, i.e. at compile time. Consequently, dynamically loading code is not necessary and even communication layout between processes may be arranged at compile time. There are examples that allow much more efficient instantiation of parallelity in a static software system (e.g. [2010hh]), but the question of how to apply these techniques to a dynamic software system remains open.

Directly addressing memory is not needed even in the context of an operating system, as the intended target architecture provides peripheral access through dedicated operations instead of memory mapped control registers, see [2016os].

Literature

[1840ob] Otto Böhtlingk (Ed.): "Pāṇini's Grammatik", Bonn 1840, Delhi 2001, ISBN 81-208-1025-2

[1945jn] John von Neumann: "First Draft of a Report on the EDVAC", Moore School of Electrical Engineering, University of Pennsylvania, June 30, 1945

[1959jb] J. W. Backus: "The syntax and semantics of the proposed international algebraic language of the Zürich ACM-GAMM conference", Proc. Internat. Conf. Inf. Proc., UNESCO, Paris, June 1959

[5] as defined for Pascal

[1963nb] J. W. Backus, F. L. Bauer, J. Green, C. Katz, J. McCarthy, A. J. Perlis, H. Rutishauser, K. Samelson, B. Vauquois, J. H. Wegstein, A. van Wijngaarden, M. Woodger, P. Naur (Ed.): "Revised Report on the algorithmic language ALGOL 60", Numerische Mathematik 4, p. 420-453 (1963)

[1975ed] Edsger W. Dijkstra: "Guarded commands, nondeterminacy, and formal derivation of programs", Comm. ACM 18, 8 (Aug. 1975), p. 453-457, EWD472

[1975jw] Kathleen Jensen, Niklaus Wirth: "Pascal: user manual and report", Springer-Verlag, 1975

[1977nw] Niklaus Wirth: "Compilerbau", Teubner, Stuttgart, ISBN 3-519-02338-5, 1977

[1978ch] C. A. R. Hoare: "Communicating sequential processes", The Queen's University, Belfast, Northern Ireland Commun. ACM 21, 8 (Aug. 1978), p. 666-677

[1978kr] Brian W. Kernighan, Dennis M. Ritchie: "The C Programming Language (1st ed.)", February 1978, Englewood Cliffs, NJ: Prentice Hall, ISBN 0-13-110163-3

[1997dw] Jochen Demuth, Stephan Weber, Sönke Kannapinn, Mario Kröplin: "Echte Compilergenerierung – Effiziente Implementierung einer abgeschlossenen Theorie", Forschungsbericht des Fachbereichs Informatik 1997/6, Technische Universität Berlin

[2001dl] Daan Leijen: "Division and Modulus for Computer Scientists", University of Utrecht, Dept. of Computer Science, December 3, 2001

[2003fy] F. Yergeau: "Request for Comments: 3629", Network Working Group, Alis Technologies, November 2003

[2007ir] ITU-T Recommendation Z.100: "Specification and Description Language (SDL)", International Telecommunication Union, 11/2007

[2009uc] The Unicode Consortium: "The Unicode Standard, Version 5.2.0", Mountain View, CA, 2009, ISBN 978-1-936213-00-9

[2010hh] James Hanlon, Simon J. Hollis: "Dynamic generation of parallel computations", Department of Computer Science, University of Bristol, UK, in proc. UK Electronics Forum 2010, (pp. 7-17), 2010

[2016os] Oskar Schirmer: "NOP – A Simple Experimental Processor for Parallel Deployment", Göttingen, 2016